

On the Robustness of Color Texture Descriptors across Illuminants

Simone Bianco¹, Claudio Cusano², Paolo Napoletano¹,
and Raimondo Schettini¹

¹ DISCo (Dipartimento di Informatica, Sistemistica e Comunicazione)
Università degli Studi di Milano-Bicocca, Viale Sarca 336, 20126 Milano, Italy
`{bianco,napoletano,schettini}@disco.unimib.it`

² Università degli Studi di Pavia, via Ferrata 1, 27100 Pavia, Italy
`claudio.cusano@unipv.it`

Abstract. In this paper we evaluate several extensions of Local Binary Patterns to color images. In particular, we investigate their robustness with respect to changes in the illuminant color temperature. To do so, we recovered the spectral reflectances of 1360 texture images from the Outex 13 data set. Then, we rendered the images as if they were taken under 33 different illuminants. For each combination of a training and test illuminant, we measured the classification performance of the texture features considered. The results of this extensive experimentation are reported and critically discussed.

Keywords: Color texture classification, illuminant invariance, reflectance recovery, Local Binary Patterns, Outex texture database.

1 Introduction

A wide variety of texture features has been proposed in the literature; see [11] for a review on this subject. Among these, Local Binary Patterns (LBP) [16] became very popular because of their simplicity, flexibility, and good discriminative power, as demonstrated by the results obtained in dealing with several application domains [14,1,12].

As common in the texture analysis field, LBPs were first designed for working with gray-scale images [15]. However, it has been proven that color can be very important for image analysis, and especially for the analysis of natural images [8]. Therefore, several approaches have been investigated to extend LBPs in order to take into account color information as well [3,17,19]. One of the shortcomings of color-aware features is that they can result overly sensitive to undesired chromatic variations, such as those introduced by changes in lighting conditions.

Mäenpää and Pietikäinen [9] reported an extensive evaluation of gray-level and color features for texture classification. Their evaluation included also several variants of LBPs. Their results seem to suggest that, gray-level LBPs outperform any other gray and color texture feature considered. In particular, the performance reported were very poor for color texture features under varying

lighting conditions. The authors concluded that color information is not only useless, but even misleading for texture classification. The significance of these findings is limited by the fact that the data sets considered contain at most images acquired under three different illuminants. In the experiments reported, one illuminant was used for the training images, and two for the testing images (in practice only two illuminant variations were considered of about 500 and 1200 Kelvin degrees respectively).

In this work we have conducted a thorough experimentation of several LBPs variants under systematic changes of the illuminant with the aim of better understanding the influence of lighting variations in color texture classification. The systematic variations of the lighting conditions have been simulated by first recovering the spectral reflectances of each texture image [2] and then using the color formation equations to render the recovered spectral image under the chosen illuminants. This approach have allowed us to investigate both small and large variations in the illuminant color (up to 9000 Kelvin degrees). Moreover, it also allowed us to consider only the variations in the illuminant color, disregarding the actual position of the light source.

The results we have obtained demonstrate that: i) the use of color information is very effective for texture classification when the color temperature of the illuminant is the same (or very similar) for training and test sets; ii) color-based texture features are not good enough to deal with large changes in color temperature. A comparison of the results obtained on actual and simulated images also suggests that the position of the light source may greatly affect the texture appearance and consequently the classification performance.

2 Color Texture Descriptors Based on Local Binary Patterns

Despite their remarkable discriminative power, the definition of Local Binary Patterns (LBP) is quite simple. Their robustness comes from the invariance with respect to rotations and contrast. Moreover, the descriptor is very compact, and can be computed very quickly. Briefly, LBPs are computed by thresholding the intensity of the neighbors of each pixel with the intensity of the pixel itself (see Figure 1). More in detail, given the number n of neighbors and the neighborhood radius r , for each pixel the $LBP_{n,r}$ code is computed as follows:

$$LBP_{n,r}(g_c) = \sum_{p=0}^{n-1} s(g_p - g_c)2^p, \quad (1)$$

where g_c is the gray level of the current pixel, g_0, \dots, g_{n-1} are the gray levels of its neighbors, and s is the step function defined as $s(x) = 1$ if $x \geq 0$, $s(x) = 0$ otherwise. The n neighbors considered lie on a circular neighbor, of radius r , of the current pixel: the gray value g_p is obtained by interpolating the intensity image at a displacement $(r \cos(2\pi p/n), r \sin(2\pi p/n))$.




Fig. 1. Local Binary Patter extraction: for each pixel, a circular neighborhood is considered; a string of binary digits is obtained comparing each neighbor against the value of the central pixel

Rotation invariance can be achieved by rotating the pattern n times, and by taking the one corresponding to the smallest integer. In terms of image pixels this corresponds to rotating the neighbor set so many times that the longest sequence of zeros is placed in front of the code.

In practice, not all the 2^n patterns are equally significant. Usually, only patterns describing a somewhat regular neighbor are considered. These patterns are called “uniform” and are defined as those for which there are at most two 0/1 transitions between adjacent bits in the code.

For classification purposes, LBPs (rotation invariant or not) are usually aggregated to form a histogram of their occurrences. Note that in the case of uniform LBPs, all non uniform patterns are accounted for in a single bin of the histogram. Typical measures of histogram similarity (such as histogram intersection and χ^2 similarity) are then used to compare pair of features.

2.1 Local Binary Patterns for Color Images

LBP features can be extended in several ways to take into account color information. The easiest one is to compute them on each color channel independently, and to concatenate the obtained histograms into a single feature vector [9]. This approach can be directly applied to the Red, Green and Blue channels of the input RGB image, or to the components of any other color space (HSV, Ohta. . .). Within this approach, co-occurrences of local patterns across different color channels can not be modeled.

Opponent Color LBP (OCLBP) [3] tries to incorporate cross-channel information in the LBP framework: LBPs are extracted from each color channel independently, and then for pairs of color channels so that the center pixel is taken from one channel and the neighboring pixels from the other. In total, six histograms are computed: three for the R, G and B channels, and three for the combinations R/G, R/B and G/B. The histograms are then concatenated to form the feature vector.

A different approach has been proposed by Porebski *et al.* [17] who considered the vectorial differences between the central pixel and its neighbors. Vectors are ordered according to a partial color order relation based on Euclidean distances. The result of the comparisons is used to populate a single color LBP histogram.

Connah and Finlayson [4] used 3D joint histograms of LBP values computed from the three R,G,B channels. In order to reduce the number of combinations of different patterns on the three channels, the patterns considered are restricted to the $LBP_{8,1}$ operator with rotation invariance and uniform patterns. Therefore, their final histogram has $10 \times 10 \times 10 = 1000$ components.

3 Artificial Illuminant Change

In order to investigate the robustness of color texture features under varying illuminants we have designed a computational procedure to synthesize color texture images as acquired under different illuminants ranging from about 3000 to 12000 Kelvin. Given a texture image, the spectral reflectance of each pixel is first recovered and then rendered under the chosen illuminant.

3.1 Reflectance Recovery

Given the CIE XYZ tristimulus values of a reflectance $r(\lambda)$ under a known illuminant $I(\lambda)$, the reflectance recovery amounts to invert the color formation equation:

$$XYZ = k \int_{\omega} r(\lambda)I(\lambda)\mathbf{C}(\lambda)d\lambda \tag{2}$$

with

$$k^{-1} = \int_{\omega} I(\lambda)\bar{y}(\lambda)d\lambda \tag{3}$$

where ω is the wavelength range of the visible light spectrum, and $\mathbf{C}(\lambda) = [\bar{x}, \bar{y}, \bar{z}]$ are the CIE XYZ color matching functions.

Different spectra recovery techniques exist [6,10,20,2]. In this work we have used the method described in [2] that exploits an adaptive linear model to synthetically describe reflectances. Given a set of object reflectances (the training set), a dimensionality reduction technique is used to derive a linear 3–dimensional basis $\mathbf{V}_0(\lambda)$ for it. The spectra in the training set are then recovered using the basis $\mathbf{V}_0(\lambda)$ and the spectral residuals are computed. The same dimensionality reduction technique is used to derive a second linear 3–dimensional basis $\mathbf{B}(\lambda)$ for them.

For each new reflectance that we want to recover, the basis $\mathbf{V}_0(\lambda)$ is used to give an initial estimate $r_0(\lambda)$ of it. The likelihood between the initial estimate $r_0(\lambda)$ and the reflectances in the training set is then measured using the Goodness of Fit Coefficient (GFC). The reflectances in the training set with a normalized likelihood above a predefined threshold are selected and a new, adaptive, linear 3–dimensional basis $\mathbf{V}_1(\lambda)$ is derived for them. If there are not enough elements to derive $\mathbf{V}_1(\lambda)$, $\mathbf{V}_1(\lambda)$ is set equal to $\mathbf{V}_0(\lambda)$.

The six coefficients $a_i, i = 1, \dots, 6$ to compute the final reflectance $\hat{r}(\lambda)$ using the two basis $\mathbf{V}_1(\lambda)$ and $\mathbf{B}(\lambda)$, i.e.

$$\hat{r}(\lambda) = \mathbf{V}_1(\lambda) [a_1 \ a_2 \ a_3]^\top + \mathbf{B}(\lambda) [a_4 \ a_5 \ a_6]^\top \tag{4}$$

are found solving the following optimization problem:

$$[a_1 \cdots a_6]^\top = \min_{x \in \mathbf{R}^6} \alpha \Delta E_{94} + \beta u_+(\hat{r}(\lambda)) + \gamma u_-(\hat{r}(\lambda)) + \delta(1 - GFC_{max}) \quad (5)$$

where the ΔE_{94} is the colorimetric error between the XYZ tristimulus values of the reflectance that we want to recover and the ones of the candidate reflectance; $u_+(\cdot)$ is a penalty term for the parts of the candidate reflectance above 1; $u_-(\cdot)$ is a penalty term for the parts of the candidate reflectance below 0; GFC_{max} is the maximum likelihood among the candidate reflectance and the reflectances in the training set. The weights in Eq. 5 are set as $[\alpha, \beta, \gamma, \delta] = [1, 100, 100, 1]$ according to [2].

3.2 Image Rendering

Once the reflectance has been estimated for every pixel of each texture, the resulting multispectral images have been rendered under 33 different illuminants belonging to four different groups of illuminant types: 8 blackbody radiators (from 3000K to 10000K with step of 1000K); 9 CIE Daylight illuminants (from D40 to D120, i.e. from 4000K to 12000K with step of 1000K); 4 CIE standard illuminants (i.e. A, B, C, and E); 12 CIE Fluorescent illuminants (from F1 to F12). Figure 2 shows the normalized Spectral Power Distribution (SPD) of the illuminants with the corresponding correlated color temperatures.

Given the recovered reflectance $r(x, y, \lambda)$ at a given pixel with coordinates (x, y) , the illuminant SPD $I(\lambda)$ under which we want to render the image, and the camera spectral transmittances $C_i(\lambda)$, $i = R, G, B$ the rendered image value at the pixel with coordinates (x, y) can be obtained using the color formation equation reported in Eq.2 using k as a normalization factor to obtain a maximum unitary sensor response, i.e.:

$$k^{-1} = \max_{x, y, i} \int_{\omega} r(x, y, \lambda) I(\lambda) C_i(\lambda) d\lambda . \quad (6)$$

4 Experimental Setup

To understand the role of lighting variations in color texture classification we firstly set up an image database by artificially changing the illumination conditions of the test suite Outex 13 and then we experimented several LBP descriptors on this database.

Outex 13 is part of the Outex collection¹ and contains 1360 images of 68 different classes [13]. The images have been obtained by subdividing 68 photographs, at 100 dpi, in 20 sub-images of size 128×128 each. The photographs depict the textures of various materials, and have been obtained under controlled

¹ Available on the web site <http://www.outex.oulu.fi/>

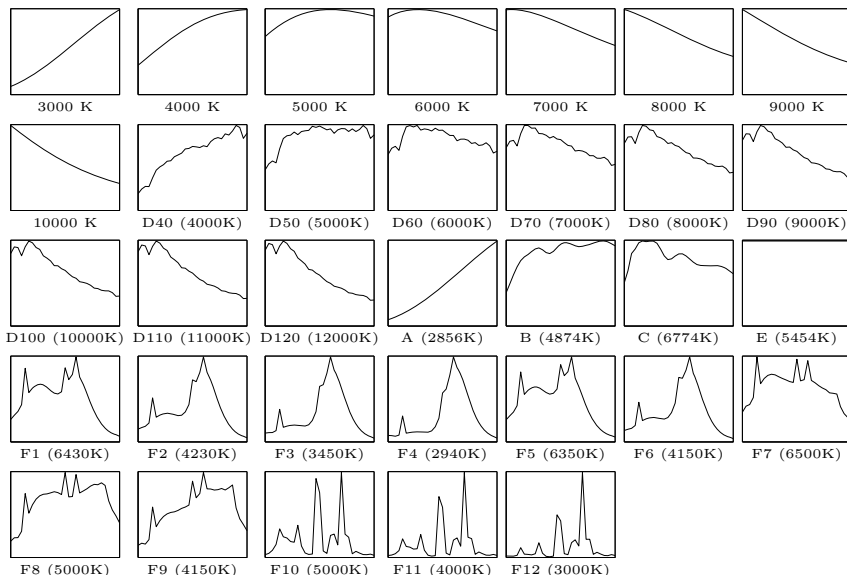


Fig. 2. Spectral power distribution of all simulated illuminants. The x axis represents the wavelength and y axis the normalized SPD.

light conditions, namely the 2856 K incandescent CIE A denoted as ‘inca’. The photographs have been acquired with a Sony DXC-755P camera, which has been calibrated using the ‘inca’ illuminant. Figure 3 shows an example for each of the 68 classes.

In order to render the Outex 13 textures under different illuminants we first recovered the spectral reflectances of each texture. Then, for each illuminant, we used the color formation equation with the Sony DXC-755P sensor spectral transmittances (which are given together with the Outex dataset) to obtain the rendered image.

Note that in order to apply the chosen spectra recovery technique [2], the Outex 13 textures are first converted into the CIE XYZ color space using the transformation derived in [9]. The training set of spectral reflectances used for the spectral recovery technique is composed of 170 measured reflectances of natural objects [18]. Figure 4 shows the rendered images for an example of Outex 13.

The Outex database contains also test suites obtained by applying to the same textures variations of scale, rotation and lighting conditions. Outex 14 is the most used test suite for evaluating the robustness of texture descriptors to illumination variations [9]. This set of images contains both the textures of Outex 13 and two variations of the same textures obtained using two different illuminants each differently positioned with respect to ‘inca’: the 2300 K horizon sunlight (denoted as ‘horizon’) and the 4000 K fluorescent TL84 (denoted as ‘TL84’). In the experimental results section we will show results obtained by computing the LBP descriptors also on this database.




Fig. 3. The 68 classes of the Outex 13 test suite. The images are acquired under the ‘inca’ illuminant with a Sony DXC-755P camera, which has been calibrated using the ‘inca’ illuminant.

In this work, we experimented several variants of LBPs. Specifically, we considered LBP applied to the luminance Y channel, to the RGB, Lab, HSV and Ohta ($I_1I_2I_3$) [9] color spaces, to Opponent Color channels [3] (denoted as OCLBP), and we also used a variant, denoted as Color Indexing, that considers 3D joint histograms of LBP values computed from the three RGB channels [4]. In all variants, apart from the Color Indexing, we considered LBPs with rotation invariant and uniform patterns, 16 neighbors, and radius 2.

In all the experiments we used the Nearest Neighbor classifier: given a texture image, its distance with respect to all the training images is computed. The prediction of the classifier is the class of the closest training set image. Performance are reported as classification rates (i.e. the ratio between the number of correctly classified images and the number of test images).

5 Experimental Results

In the first experiment we have studied how much the classification performance is affected by the differences in the illuminant color between the training and test sets. Since we considered 33 illuminants, we have $33 \times 33 = 1089$ classification experiments, one for each combination of training and test illuminants. Figure 5 shows the results obtained by using the seven variants of LBP. The results are reported in terms of classification rate as function of the absolute difference (ΔT) between the temperature of the training and test illuminants.

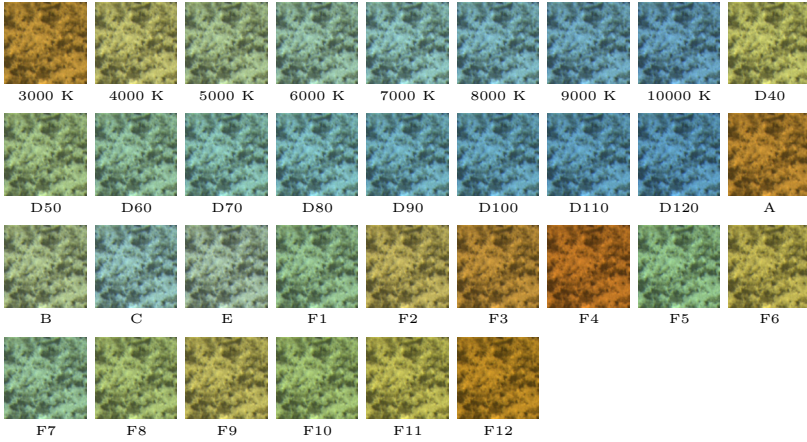


Fig. 4. Examples of a Outex 13 class rendered under the 33 different illuminants using the Sony DXC-755P sensor spectral transmittances




Fig. 5. Classification rates obtained by using the seven features considered as functions of the absolute difference (ΔT) between the temperature of the training and test illuminants

The figure shows that when the illuminant is the same (i.e. $\Delta T = 0$) the use of color information outperforms LBP on the luminance image. In particular, the OCLBP method obtained the best result (about 91.5%).

With variable illuminant, the classification rates are decreasing with respect to ΔT . For $\Delta T > 0$ the best strategy is to use LBPs in the RGB color space. This fact can be explained by considering that changes in illumination are often approximated by diagonal transformations in the RGB color space, with positive gains [7]. The order between pairs of samples is not affected by this kind of transformations. Since LBPs depends solely on these orderings, LBP-RGB

Table 1. Classification performance of all the descriptors on Outex 14

	inca	TL84	horiz		inca	TL84	horiz		inca	TL84	horiz
inca	81.0	68.5	68.2	inca	86.3	71.3	48.8	inca	88.1	43.5	58.5
TL84	71.6	79.0	77.9	TL84	72.4	86.0	46.9	TL84	45.4	88.5	44.9
horiz	68.8	76.2	82.7	horiz	58.2	57.4	88.2	horiz	67.1	47.5	88.2
(a) Y			(b) RGB			(c) HSV					
	inca	TL84	horiz		inca	TL84	horiz		inca	TL84	horiz
inca	88.5	42.4	77.9	inca	86.3	53.5	57.9	inca	93.1	8.5	15.3
TL84	47.7	84.9	40.7	TL84	54.7	83.4	54.7	TL84	8.2	96.2	2.7
horiz	77.5	43.2	89.3	horiz	66.8	54.6	89.4	horiz	23.8	2.8	90.7
(d) Lab			(e) $I_1 I_2 I_3$			(f) Opponent Color					
	inca	TL84	horiz		inca	TL84	horiz				
inca	73.8	42.8	25.7	inca	73.8	42.8	25.7				
TL84	42.1	75.7	19.7	TL84	42.1	75.7	19.7				
horiz	32.5	28.4	73.1	horiz	32.5	28.4	73.1				
(g) Color Indexing											

is approximately invariant to changes in illumination. The same holds for the luminance image. In fact, the performance of LBP-Y are constant across all the values of ΔT .

For other color spaces, such as HSV, Lab and $I_1 I_2 I_3$, where the illuminant change can not be approximated as a diagonal transformation, there can be inversions between the values in the color channels. In fact, for these three color spaces there is a noticeable drop in performance as ΔT increases (see Figure 5).

The two methods exploiting cross-channel information (OCLBP and Color Indexing) are the most affected by the changes in illumination. OCLBP shows a convex behavior: its worst result is ΔT of about 5000K. This fact can be explained by the fact that OCLBP combines six LBP histograms: three on the same channels (R/R, G/G, B/B) that should behave like LBP-RGB, and three across channels (R/G, R/B, G/B) that cause the low classification rate. These last three histograms tend to be irrelevant for large ΔT s, where the red channel (for high temperatures) and the blue channel (for low temperatures) become negligible (this is why the classification rate improves for large values of ΔT).

We compared the results obtained with artificial changes in illuminations against those obtained on the Outex-14 data set. Table 1 (a–g) report the results obtained by rotating the training and the test set across the three available illuminants. The elements on the diagonal correspond to the cases where training and test images have been taken under the same illuminant; the classification rates, here, are very similar to those obtained with the artificial illuminants when $\Delta T = 0$. The results do not agree when training and test illuminants are different (out-of-diagonal elements of the tables). In this case, the results with real illuminants are lower than those with artificial illuminants. A possible

explanation is that in the Outex-14 data set the three illuminants were placed in different positions, while in our simulations the only variability was in the illuminant color. The changes in the illuminant position may affect LBPs even more than the changes in the illuminant color. For instance, with the artificial illuminants the performance of LBP-RGB and LBP-Y are constant, while with the real data the change of illuminant caused a clear drop in performance.

6 Conclusions

In this paper we investigated the robustness of several color variants of LBPs with respect to changes in the illuminant color. To do so, we recovered the spectral reflectances of the Outex 13 images, and we rendered them as if they were taken under 33 different illuminants. For each combination of a training and a test illuminant, we measured the classification performance of the features considered. The experiments show that color information can be very useful for texture classification when there ain't any variation between training and test illuminant. They also show that some descriptors (such as LBP-Y and LBP-RGB) are almost invariant with respect to changes in the color of the illuminant. This fact does not agree with the results found in the literature. We speculate that the difference can be explained by the fact that, for real images, the changes in the illuminant color are usually paired with changes in its position. This point deserves to be investigated in greater detail by conducting additional experiments. We plan to do so in our future work; we also plan to use the insights provided by our experiments to design new descriptors for texture classification under varying lighting conditions (a first result of this line of research is reported in [5]).

References

1. Ahonen, T., Hadid, A., Pietikäinen, M.: Face description with local binary patterns: Application to face recognition. *IEEE Trans. on PAMI* 28(12), 2037–2041 (2006)
2. Bianco, S.: Reflectance spectra recovery from tristimulus values by adaptive estimation with metameric shape correction. *J. Opt. Soc. Am. A* 27, 1868–1877 (2010)
3. Chan, C.H., Kittler, J., Messer, K.: Multispectral local binary pattern histogram for component-based color face verification. In: *First IEEE Intl. Conf. on Biometrics: Theory, Applications, and Systems*, pp. 1–7 (2007)
4. Connah, D., Finlayson, G.: Using local binary pattern operators for colour constant image indexing. In: *Proc. European Conf. on Color in Graphics, Imaging, and Vision*, p. 5 (2006)
5. Cusano, C., Napoletano, P., Schettini, R.: Illuminant invariant descriptors for color texture classification. In: Tominaga, S., Schettini, R., Trémeau, A. (eds.) *CCIW 2013. LNCS*, vol. 7786, pp. 239–249. Springer, Heidelberg (2013)
6. Dupont, D.: Study of the reconstruction of reflectance curves based on tristimulus values: comparison of methods of optimization. *Color Research and Application* 27, 88–99 (2002)
7. Finlayson, G.D., Drew, M.S., Funt, B.V.: Color constancy: Generalized diagonal transforms suffice. *J. Opt. Soc. Am. A* 11, 3011–3020 (1994)

8. Haindl, M., Filip, J.: *Visual Texture*, vol. XXXI. Springer, Heidelberg (2013)
9. Mäenpää, T., Pietikäinen, M.: Classification with color and texture: jointly or separately? *Pattern Recognition* 37(8), 1629–1640 (2004)
10. Mansouri, A., Sliwa, T., Hardeberg, J., Voisin, Y.: An adaptive-pca algorithm for reflectance estimation from color images. In: 19th Intl. Conf. on Pattern Recognition, pp. 1–4 (2008)
11. Mirmehdi, M., Xie, X., Suri, J.: *Handbook of Texture Analysis*. Imperial College Press, London (2008)
12. Nanni, L., Lumini, A., Brahmam, S.: Local binary patterns variants as texture descriptors for medical image analysis. *Artificial Intelligence in Medicine* 49(2), 117–125 (2010)
13. Ojala, T., Mäenpää, T., Pietikäinen, M., Viertola, J., Kyllönen, J., Huovinen, S.: Outex-new framework for empirical evaluation of texture analysis algorithms. In: 16th Intl. Conf. on Pattern Recognition, vol. 1, pp. 701–706 (2002)
14. Ojala, T., Pietikäinen, M.: Unsupervised texture segmentation using feature distributions. *Pattern Recognition* 32(3), 477–486 (1999)
15. Ojala, T., Pietikäinen, M., Mäenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 24(7), 971–987 (2002)
16. Pietikäinen, M., Hadid, A., Zhao, G., Ahonen, T.: Local binary patterns for still images. In: *Computer Vision Using Local Binary Patterns, Computational Imaging and Vision*, vol. 40, pp. 13–47. Springer, London (2011)
17. Porebski, A., Vandenbroucke, N., Macaire, L.: Haralick feature extraction from lbp images for color texture classification. In: *First Workshops on Image Processing Theory, Tools and Applications*, pp. 1–8 (2008)
18. Vhrel, M., Gershon, R., Iwan, L.: Measurement and analysis of object reflectance spectra. *Color Research and Application*, 4–9 (1994)
19. Zhu, C., Bichot, C.E.: Multi-scale color local binary patterns for visual object classes recognition. In: *ICB 2007*, pp. 3065–3068 (2010)
20. Zuffi, S., Santini, S., Schettini, R.: From color sensor space to feasible reflectance spectra. *IEEE Trans. on Signal Processing* 56, 518–531 (2008)